

AE

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 199 851 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
24.04.2002 Bulletin 2002/17

(51) Int Cl.7: H04L 12/56, H04L 12/44

(21) Application number: 01440254.9

(22) Date of filing: 06.08.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Hill, Rex A.
San Diego, CA 92129 (US)
• Cinco, Dante
San Diego, CA 92126 (US)

(30) Priority: 20.09.2000 US 234028 P
21.11.2000 US 721101

(74) Representative: Wörz, Volker Alfred
Dreiss, Fuhlendorf,
Steimle & Becker
Patentanwälte
Gerokstrasse 6
70188 Stuttgart (DE)

(71) Applicant: ALCATEL
75008 Paris (FR)

(54) QoS shaping/provisioning for data communication switch

(57) A method and apparatus provides QoS shaping/ provisioning scheme for a data communications switch, such as a DiffServ aware router or a 802.1Q aware bridge, in which distinct internal and outbound priority values are assigned to a packet based on flow properties associated with an inbound packet. The flow properties used to assign the internal and outbound priority values may include at least one value from a packet

field that is not dedicated to defining QoS. The internal priority value provides a priority to the packet during processing in the switch. The outbound priority value is applied to the packet in lieu of the inbound priority value prior to transmitting the packet from the switch. The flow properties used to determine the internal and outbound priority values may include, for example, Layer 2, Layer 3, and Layer 4 information encoded in the packet.




FIG. 3

EP 1 199 851 A1

Best Available Copy

Description

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of the filing date of U.S. provisional patent application (Application No. 60/234,028), filed September 20, 2000 and entitled "High Speed LAN Switching Controller," the contents of which are hereby incorporated by reference. This application contains subject matter related to U.S. patent application (Attorney Docket No. 39983/JEJ/X2/134019) filed November 21, 2000 and entitled "Stage-Implemented QoS Shaping for Data Communication Switch."

FIELD OF THE INVENTION

[0002] The present invention relates to devices for determining the Quality of Service (QoS) in a data communications switch, and especially to devices for determining the internal and outbound QoS based on a number of flow properties.

BACKGROUND OF THE INVENTION

[0003] IEEE (Institute of Electrical and Electronics Engineers) Standard 802.1Q entitled "Virtual Bridged Local Area Networks" defines an industry standard for virtual bridged local area networks (VLANs). The 802.1Q standard, among other things, defines a convention for adding a tag header to a Layer 2 data packet, i.e. a "frame", in the creation of an 802.1Q-compliant packet. The tag header may include, among other things, a VLAN Identifier and a priority value assigned to the packet. The VLAN Identifier typically determines what LAN devices are authorized to receive the packet, and the priority value typically determines how fast the packet will be received by the authorized LAN devices relative to other packets.

[0004] As an 802.1Q-compliant packet passes through each bridge in a bridged LAN, the priority value in the inbound packet may be "regenerated" by the bridge to determine a corresponding priority value for the outbound packet. According to the standard, the outbound priority value may be independently determined on each bridge as a function of the port on which the bridge receives the packet and the inbound priority value. There is no provision in the standard, however, to take into account values from fields in the packet that are not dedicated to defining quality of service (QoS) when assigning the outbound priority value.

[0005] Therefore, it is desirable to take into account values from fields in the packet that are not dedicated to defining QoS when assigning the outbound priority value. It may also be desirable to assign distinct internal and outbound priority values to the packet, the former for application in prioritizing the packet on the current bridge and the latter for application to the packet as transmitted from the current bridge.

SUMMARY OF THE INVENTION

[0006] In one embodiment of the present invention, a quality of service (QoS) shaping and provisioning method for a switch having a plurality of ports is provided. The method includes the steps of receiving a packet having a first priority value on a first port, determining a second priority value for the packet based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS, and processing the packet based on the second priority value.

[0007] In another embodiment of the present invention, a QoS shaping and provisioning method for a switch having a plurality of ports is provided. The method includes the steps of receiving a packet having a first priority value on a first port, determining a second priority value for the packet based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS, and applying the second priority value to the packet prior to transmission from the switch.

[0008] In yet another method of the present invention, a QoS shaping and provisioning method for a switch having a plurality of ports is provided. The method includes the steps of receiving a packet having a first priority value on a first port, determining a second priority value for the packet based on one or more flow properties, processing the packet based on the second priority value, determining a third priority value for the packet based on the one or more flow properties, and applying the third priority value to the packet prior to transmission from the switch.

[0009] In yet another embodiment of the present invention, a switch having one or more switching modules that are capable of QoS shaping and provisioning is provided. Each switching module includes one or more ports for receiving a plurality of inbound packets and for transmitting a plurality of outbound packets, an access controller coupled to the input ports for receiving the inbound packets, each inbound packet having an inbound priority value and a plurality of flow properties, and a switching controller coupled to the access controller for receiving the inbound packets from the access controller, for generating one or more packet priority values based on the plurality of flow properties, and for providing the outbound packets to the ports to be transmitted.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] These and other aspects of the invention may be understood by reference to the following detailed description, taken in conjunction with the accompanying drawings, which are briefly described below.

FIG. 1 is a switch in one embodiment of the present invention;

FIG. 2 is a switching module, which may be used to

in the switch of FIG. 1;

FIG. 3 illustrates mapping of inbound priority values into outbound priority values in one embodiment of the present invention;

FIG. 4 is QoS shaping maps in one embodiment of the present invention;

FIG. 5 is a flow diagram describing the QoS shaping/provisioning in accordance with one embodiment of the present invention; and

FIG. 6 is a flow diagram describing stage-implemented QoS shaping in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

[0011] One embodiment of the present invention provides a QoS shaping/provisioning scheme for a switch, such as 802.1Q aware bridge. The QoS shaping/provisioning scheme broadly includes many aspects of QoS-related tasks, and may include but not limited to priority shaping, link sharing, bandwidth provisioning and bandwidth limiting. The switch may also be referred to as a data communication switch.

[0012] An analogous QoS shaping/provisioning scheme may also be applied to other types of switches, such as DiffServ (differentiated services) aware routers. In DiffServ aware routers, the QoS shaping/provisioning scheme may be implemented with respect to a DiffServ field in the header of a Layer 3 packet, e.g. IP datagram, as specified in IETF Request for Comment 2475 entitled, "An Architecture for Differentiated Services" (RFC 2475).

[0013] In one embodiment of the present invention, distinct internal and outbound QoS values preferably are assigned by a switch to an inbound packet based on a plurality of flow properties associated with the inbound packet.

[0014] QoS is typically defined by fields, such as, for example, (1) the 3-bit 802.1Q tag field in Layer 2, (2) the 8-bit Type of Service (ToS) field in Layer 3, and (3) the 6-bit DiffServ in Layer 3 that are dedicated to defining QoS.

[0015] The flow properties preferably also include, but are not limited to, at least one value from a packet field that is not dedicated to defining QoS. For example, the plurality of flow properties used to determine the internal and outbound QoS values may include Data Link Layer (Layer 2), Network layer (Layer 3) and Transport Layer (Layer 4) information in the packet at the time of receiving the inbound packet on the switch as well as other values that may be derived from Layer 2, Layer 3 and Layer 4 information during processing of the packet.

[0016] For example, the values from packet fields that are not dedicated to defining QoS, which are used to determine internal and outbound QoS values, may include but are not limited to source and destination addresses such as Layer 2 and Layer 3 addresses as well as Layer 4 port numbers, which may also be referred to

as socket IDs. Other flow properties that are not dedicated to defining QoS but are used to determine internal and outbound QoS values may include values not necessarily in the packet header but are derived during the classification process, such as, for example, the physical port number on which the packet arrived and the VLAN Identifier onto which the packet is classified.

[0017] For example, for a server and a PC connected to a switch through the same port and requesting the same priority, it may be desirable to respect the server's requested priority while overriding the PC's requested priority. On the other hand, it may be desirable to respect both the server's and PC's requested priority when transmitting the packet from the current switch but to override the PC's requested priority when prioritizing the packet on the current switch. The switch may distinguish between packets from the server and the PC based on values such as Layer 2 and Layer 3 addresses.

[0018] The internal QoS value preferably is applied to provide the specified QoS to the inbound packet while being processed in the switch, whereas the outbound QoS value preferably is applied to the inbound packet in lieu of the inbound QoS value prior to transmitting the inbound packet from the switch as an outbound packet. For example, packets having higher internal priorities may receive expedited processing, expedited queuing and/or preferential bandwidth treatment.

[0019] The inbound QoS value, the internal QoS value and the outbound QoS value may, for example, also be referred to as an inbound priority value, an internal priority value and an outbound priority value, respectively. The inbound QoS value and the outbound QoS value may also be referred to as an ingress priority value and an egress priority value, respectively, or as an ingress QoS value and an egress QoS value, respectively. The inbound packet and the outbound packet may, for example, also be referred to as an ingress packet and an egress packet, respectively.

[0020] FIG. 1 illustrates a switch 100 in one embodiment of the present invention. Switch 100 preferably includes a switching backplane 111 coupled to switching modules 101a-101c. The switching backplane 111 may include a switching fabric, or switching fabric functions may be distributed between the switching backplane and the switching modules. The switching fabric (not shown) may also interface between the switching backplane 111 and the switching modules 101a-101c.

[0021] The switching modules 101a-101c preferably are coupled to one another over control paths 103a and 103b, respectively, and each switching module preferably is associated with one or more LANs 121a-121c. The LANs 121a-121c preferably include one or more virtual local area networks (VLANs). The switching modules communicate with the LANs over data interfaces 104a-104c.

[0022] The switching modules preferably interface with the switching backplane 111 over data paths 102a-102c, respectively. The switching modules preferably

transmit packet data to the switching backplane 111 and receive packet data from the switching backplane 111. Depending on the source and destination addresses, an inbound packet from a switching module may be provided over the switching backplane 111 as an outbound packet to the same or different switching module.

[0023] The switching modules 101a-101c and the LANs 121a-121c are shown for illustrative purposes only. There may actually be more or less switching modules and LANs in this and other embodiments of the present invention. Further, in other embodiments, the switch may have a configuration that is different from the configuration of the switch 100 of FIG. 1.

[0024] FIG. 2 is a switching module 200 in one embodiment of the present invention. The switching module 200 may be used as one of the switching modules 101a-101c of FIG. 1. The switching module 200 preferably includes a stage-implemented QoS shaping/provisioning scheme in which packet processing and QoS shaping/provisioning scheme are performed in multiple stages. The switching module 200 includes an access controller 201 coupled to one or more LANs over one or more data interfaces 202 and a switching controller 211 over an internal interface 204. The LANs may include VLANs.

[0025] The access controller 201 may include a media access controller (MAC). The access controller 201 preferably receives inbound packets from the LANs, formats them, and transmits them to the switching controller 211. The access controller 201 preferably also receives outbound packets from the switching controller 211, formats them, and transmits them to the LANs.

[0026] The switch 200 preferably also includes other elements that are used to facilitate packet processing operations. These switch elements preferably include but are not limited to a packet forwarding logic 220 and a QoS shaping logic 215. The QoS shaping logic 215 and the packet forwarding logic 220 preferably are coupled to the switching controller 211. The switch 200 may also include other sub-modules for performing various different tasks.

[0027] The switch 200 preferably is an ingress buffered switch, in which most of the large buffers are located at the inbound side. Thus, internal and outbound priority value determinations preferably are made at the inbound side where most of the buffers are located. In other embodiments, determination of the internal and outbound priority values may be done at the outbound side. In still other embodiments, determination of the internal and outbound priority values may be performed by distributed processing, in which some of the processing for determination of the priority values is done at the inbound side and some processing is done at the outbound side. If the priority values are added to the packet header, for example, QoS shaping/provisioning may be performed at the outbound side using outbound buffers.

[0028] QoS shaping may also be performed at the outbound side in order to perform true shaping. True shaping typically requires that packets wait very close

to the link on which they are going to be transmitted and typically also requires a sort operation on those packets. When there is a huge set of packets waiting around in the outbound buffers, QoS shaping may be desirable to figure out which packet to send next.

[0029] The switching controller 211 preferably receives inbound packets, preferably subjects selected ones of them to internal and outbound QoS shaping, and transmits selected ones of them to the access controller 201 over the internal interface 204. In this embodiment, QoS shaping preferably is provided with the assistance of the QoS shaping logic 215 and the packet forwarding logic 220. In other embodiments, QoS shaping may be provided using other elements and/or logic.

[0030] An inbound packet to the switching controller 211 typically has an inbound priority value in the Layer 2 802.1Q/p header or the Type of Service (ToS) field in the Layer 3 IP header. The fields including three bits from either the 802.1Q/p header or the IP header may be used, but both fields typically are not used at the same time. For example, the inbound priority may be included in a 3-bit ToS precedence (ToSp) sub-field of the ToS field. The ToS field in the IP header may also be referred to as a differentiated services code point (DSCP) field or a DiffServ field.

[0031] When the inbound packet enters the switching controller 211, the switching controller preferably provides the inbound packet header to the packet forwarding logic 220. In other embodiments, the switching controller may provide the whole inbound packet to the packet forwarding logic. The packet forwarding logic includes a forwarding information finder 222 and a forwarding index finder 224. The forwarding information finder 222 may also be referred to as a header cache, and the forwarding index finder 224 may also be referred to as a look up engine or a pseudo-cam (pseudo-content addressable memory).

[0032] The forwarding index finder 224 preferably receives the inbound packet header. The forwarding index finder 224 preferably uses the Layer 2, 3, or 4 header information to access a routing table for an index. Addressing and other information in the header that is used to look up the index is also referred as a key or as an input key. The forwarding index finder preferably classifies the packet into a flow, which has an index associated with it. The index for the flow may also be referred to as an ID or a flow ID. Each flow preferably is associated with a unique index, and the packets that are classified onto the same flow preferably are associated with the same index. Each flow preferably is also associated with a priority class, which is going to be described in detail later in reference to internal and outbound priority mappings.

[0033] The forwarding index finder 224 preferably provides the index to the switching controller 211. The switching controller 211 preferably uses the index to look up a table entry in a large table associated with the forwarding information finder 222. The large table asso-

ciated with the forwarding information finder may include, e.g., a next hop header table (NHHT).

[0034] The table entry, which may also be referred to as a flow descriptor, preferably has a number of fields for each packet on the flow. The table entry may contain information for performing tasks such as, for example, queue ID (QID) to specify the port number to be used by an outbound packet, maximum frame length, whether or not to perform 802.1Q tag stripping, and whether or not the packet will be processed by software. For example, the table entry in one embodiment preferably contains Layer 2 and Layer 3 information such as information on how to modify the header, VLAN ID, and a 6-bit priority description index (PDI), which may also be referred to as a QoS shaping index. In other embodiments, the table entry may also contain Layer 4 information. The PDI preferably facilitates mapping of inbound packets to different priority classes. In other embodiments, the number of bits in the PDI may be more or less than six.

[0035] The table entry preferably is then provided to the switching controller 211 to be used to further process the inbound packet. The switching controller preferably also provides the table entry to the QoS shaping logic 215 for packet processing including but not limited to priority shaping, link sharing, bandwidth provisioning and bandwidth limiting. The QoS shaping logic 215 includes QoS shaping maps 217 and a QoS shaper 219. The QoS shaper may include a queue manager, and functions of the QoS shaper may be distributed throughout the switching modules and the switching fabric (not shown) that interfaces between multiple switching modules. The QoS shaping maps 217 preferably receive the PDI and the 3-bit inbound priority value from the packet and preferably performs a table lookup for internal and outbound priority values.

[0036] FIG. 3 is an exemplary block diagram illustrating mapping of inbound 802.1p priority value and an inbound ToS priority value into an internal priority value, an outbound 802.1p priority value, an outbound ToS priority value, and an outbound DSCP (differentiated services code point).

[0037] The inbound priority values represented by inbound 802.1p value 300 and an inbound ToS value 302 preferably are used to derive an internal priority value 305 and an outbound priority value 307. The inbound 802.1p value is included in Layer-2 802.1Q/p header, and the ToS value is included in Layer-3 IP header. Depending on the mode of operation, the inbound priority value used may be the inbound 802.1p priority value 300 or the inbound ToS priority value 302. In addition, the outbound priority value 307 may be an outbound 802.1p priority value 312 or an outbound ToS priority value 314.

[0038] In other embodiments, the inbound 802.1p value and the inbound ToS value may also be provided to a DSCP table 308. For example, an 8-bit inbound DSCP value may be identical to the 8-bit inbound ToS value. The DSCP table 308 includes an outbound DSCP map

309 and an outbound DSCP mask 310. An outbound DSCP value 311 preferably is provided, depending on the mode of operation, either as a DSCP value in the table entry of the forwarding information finder, or as the modified inbound DSCP value. The inbound DSCP value may be modified through mapping to the DSCP map 309 and the outbound DSCP mask 310, wherein the DSCP mask may be used to modify some of the bits in the DSCP map.

[0039] FIG. 4 illustrates a set of shaping maps 400 within the QoS Shaping Maps 217 in one embodiment of the present invention. The QoS Shaping Maps 217 include a plurality of priority descriptor rows. The shaping maps 400 include internal priority maps 404 and outbound priority maps 406. Each priority descriptor row in this embodiment preferably includes an internal priority map including eight 2-bit values, and an outbound priority map including eight 3-bit values. Since the PDI is six bits in size in this embodiment, there may be up to 64 possible different priority descriptor rows that may be accessed using the PDI as an index. In other embodiments, the PDI may be more or less than six bits in size, and correspondingly, there may be more or less than 64 possible different priority descriptor rows.

[0040] Since internal priority values have only two bits and can have values between 0-3 in this embodiment, the 3-bit inbound priority values preferably are mapped to a 2-bit value during internal priority mapping. On the other hand, since outbound priority values have three bits, the inbound priority values may be mapped to the outbound priority values without losing any significant bits.

[0041] Each priority descriptor row preferably is associated with a priority class. The priority classes preferably are derived during packet classification from information in the packet header such as Layer 2 and Layer 3 source and destination addresses and Layer 4 port numbers. Each priority class preferably has the internal and outbound priority mapping values that are different from the internal and outbound priority mapping values of other priority classes. In other embodiments, two or more priority descriptor rows associated with different priority classes may have the same internal and/or outbound priority mapping values.

[0042] For example, in this embodiment, the priority classes may include PASS, FORCE and INVERT classes 403 associated with priority descriptor indices (PDIs) 402. These priority classes, for example, may be selected based on the source addresses of the inbound packets.

[0043] Each of the PDIs 402 preferably is associated with a particular priority class. For example, in FIG. 4, the priority index <000000> is associated with the PASS class, the priority index <000001> is associated with the FORCE class, and the priority index <111111> is associated with the INVERT class. Further, the PASS class may be associated with a server, the FORCE class may be associated with a PC, and the INVERT class may be

associated with an erroneously configured network device.

[0044] As the name "pass" suggests, the outbound priority map associated with the PASS class preferably passes the inbound priority value and provides as an outbound priority value. Since only two bits are available in each internal priority value in this embodiment, two most significant bits of the inbound priority values preferably are used as internal priority values.

[0045] As the name "force" suggests, the internal and outbound priority maps associated with the FORCE class preferably force the internal and outbound priority values to 3 and 7, respectively. In other embodiments, other internal and outbound priority values may be associated with the FORCE class. In still other embodiments, there may be multiple different FORCE classes having different internal and outbound priority values.

[0046] As the name "invert" suggests, the internal priority map associated with the INVERT class preferably inverts the two most significant bits of the inbound priority value to provide as an internal priority value. Further, the outbound priority map inverts the entire inbound priority value to provide as an outbound priority value.

[0047] For example in FIG. 4, an internal priority map of <0,0,1,1,2,2,3,3> and an outbound priority map of <0,1,2,3,4,5,6,7>, respectively, are provided for PDI of <000000>. Therefore, when the inbound packet with a source address having the priority class of PASS is received, it may be provided with the PDI of <000000>. If that inbound packet has an inbound priority value of 3, the inbound priority value is mapped to an internal priority value of 1 and an outbound priority value of 3. Similarly, when inbound packets with source addresses having the priority classes of FORCE and INVERT are received, they may be provided with the PDI of <000001> and <111111>, respectively. Other priority classes and other methods of mapping may be used in other embodiments.

[0048] If a new set of mapping data is desired, the user may use a QoS shaping map update interface 228 of FIG. 2 to enter new internal and/or outbound maps into the shaping maps 400 at a PDI location from <000000> to <111111>. A shaping map update signal as well as the map data for updating may be provided to the QoS shaping maps over the QoS shaping map update interface 228.

[0049] In this and other embodiments, flow properties other than the source address, a destination address for example may be used to select the type of mapping used to convert the inbound priority values to internal priority values and outbound priority values, respectively.

[0050] Referring back to Figure 2, after the internal and outbound priority values are identified, the QoS shaping maps 217 preferably return these priority values to the switching controller 211. The switching controller 211 preferably modifies the inbound packet to in-

clude the new priorities. The switching controller then sends the inbound packet having the new internal and outbound properties to the QoS shaper 219 so that the inbound packet can be put in an order for transmit as an outbound packet based on its new priority values.

[0051] The new outbound priority value is typically not used by the switching module 200. A switching module in a next switch preferably uses this new outbound priority value to generate the next set of internal and outbound priority values.

[0052] When the packet's priority is selected based on the new internal priority value, the QoS shaper provides the packet to the switching controller 211, which then provides the packet to the access controller 201 to be transmitted over the data interfaces 202 to the LANs.

[0053] FIG. 5 is a flow diagram of QoS shaping/provisioning in one embodiment of the present invention. An inbound packet having QoS is received by an access controller in a switch in step 510. The access controller, e.g., may be the access controller 201 of FIG. 2. The QoS may include but is not limited to a 802.1p priority value and a ToS priority value. If the inbound packet does not have QoS, it may be given a default priority based on its inbound port and protocol. The default priority may also be based on other flow properties such as Layer 2 and Layer 3 source and destination addresses, as well as Layer 4 port numbers.

[0054] The access controller preferably processes the inbound packet and sends it on to a switching controller, which may be similar to the switching controller 211 of FIG. 2. The switching controller in step 520 preferably determines new internal and outbound QoS for the inbound packet based on the flow properties, which may include a value from packet fields that are not dedicated to defining QoS, such as, for example, Layer 2 and Layer 3 source and destination addresses as well as Layer 4 port numbers.

[0055] The switching controller 211 in step 530 preferably applies the new internal QoS to the inbound packet to provide a QoS prioritization to the inbound packet within the switch. The switching controller in step 540 inserts the outbound QoS to the inbound packet. The inbound packet is then provided to the switching module coupled with the LAN to which the inbound packet is provided as an outbound packet. The outbound packet is transmitted with the new QoS outbound priority. Depending on source and destination addresses of the packet, the switching module that transmits the outbound packet may be the same or different from the switching module that received the inbound packet, since the inbound packets received by one switching module may be provided to another switching module during packet processing over the switching backplane.

[0056] FIG. 6 is a flow diagram of stage-implemented QoS shaping in one embodiment of the present invention. For example, step 520 of FIG. 5 may be implemented in stages as illustrated in FIG. 6. A QoS shaping index preferably is determined from the flow properties of an

inbound packet in step 610, for example, in the packet forwarding logic 220 of FIG. 2. Then, an internal and outbound QoS shaping maps, e.g., in the QoS shaping maps 217, preferably are selected based on the QoS shaping index in step 620. Further, internal and outbound QoS values preferably are selected using the inbound QoS in step 630.

[0057] It will be appreciated by those of ordinary skill in the art that the invention can be embodied in other specific forms without departing from the spirit or essential character hereof. The present description is therefore considered in all respects to be illustrative and not restrictive. The scope of the invention is indicated by the appended claims, and all changes that come within the meaning and range of equivalents thereof are intended to be embraced therein.

[0058] For example, a switch may be used that has a configuration that is different from the switch described in the specification. For example, in the described embodiments, QoS shaping preferably is performed in a QoS shaper of the switching modules. In other embodiments, QoS shaping function may be a distributed process where some or all of QoS shaping is performed in other components such as a switching fabric. Further, in the described embodiments, QoS shaping preferably is performed by the switch during inbound processing of the packets. In other embodiments, QoS shaping may be performed by the switch during outbound processing of the packets. In still other embodiments, some parts of QoS shaping/provisioning may be performed during inbound processing and other parts of QoS shaping/provisioning may be performed during outbound processing.

Claims

1. A quality of service (QoS) shaping and provisioning method for a switch having a plurality of ports, the method comprising the steps of:

receiving a packet having a first priority value on a first port;
determining a second priority value for the packet based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS; and
processing the packet based on the second priority value.

2. The method of claim 1 wherein the first priority value includes information from at least one of a 802.1Q tag field in Layer 2, a Type of Service (ToS) field in Layer 3 and a Diffserv in Layer 3.
3. The method of claim 1 wherein the first priority value includes an inbound priority value, and the second priority value is determined based on the inbound

priority value.

4. The method of claim 1 wherein the second priority value includes an internal priority value, and the internal priority value is used to indicate a processing priority of the packet while the packet is being processed in the switch.
5. The method of claim 1 wherein the second priority value is determined based on one or more flow properties including at least one of Layer 2 information, Layer 3 information and Layer 4 information.
6. The method of claim 1 wherein the first priority value is mapped by a QoS shaping map into the second priority value.
7. The method of claim 6 wherein mapping information of the QoS shaping map can be updated during operation of the switch.
8. A quality of service (QoS) shaping and provisioning method for switch having a plurality of ports, the method comprising the steps of:

receiving a packet having a first priority value on a first port;
determining a second priority value for the packet based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS; and
applying the second priority value to the packet prior to transmission from the switch.
9. The method of claim 8 wherein the first priority value includes an inbound priority value, and the second priority value is determined based on the inbound priority value.
10. The method of claim 8 wherein the second priority value includes an outbound priority value, and the outbound priority value is applied to the packet before the packet is transmitted from the switch.
11. The method of claim 8 wherein the second priority value is determined based on one or more flow properties including at least one of Layer 2 information, Layer 3 information and Layer 4 information.
12. The method of claim 8 wherein the first priority value is mapped by a QoS shaping map into the second priority value.
13. The method of claim 12 wherein mapping information of the QoS shaping map can be updated during operation of the switch.
14. A quality of service (QoS) shaping and provisioning

method for a switch having a plurality of ports, the method comprising the steps of:

receiving a packet having a first priority value on a first port;
determining a second priority value for the packet based on one or more flow properties;
processing the packet based on the second priority value;
determining a third priority value for the packet based on the one or more flow properties; and
applying the third priority value to the packet prior to transmission from the switch.

15. The method of claim 14 wherein the first priority value includes an inbound priority value, and at least one of the second priority value and the third priority value is determined based on the inbound priority value.

16. The method of claim 14 wherein the second priority value includes an internal priority value, and the internal priority value is used to indicate a processing priority of the packet while the packet is being processed in the switch.

17. The method of claim 14 wherein the third priority value includes an outbound priority value, and the outbound priority value is applied to the packet before the packet is transmitted from the switch.

18. The method of claim 14 wherein the second priority value is determined based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS.

19. The method of claim 14 wherein the third priority value is determined based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS.

20. The method of claim 14 wherein the second priority value is determined based on one or more flow properties including at least one of Layer 2 information, Layer 3 information and Layer 4 information.

21. The method of claim 14 wherein the third priority value is determined based on one or more flow properties including at least one of Layer 2 information, Layer 3 information and Layer 4 information.

22. The method of claim 14 wherein the first priority value is mapped by a first QoS shaping map into the second priority value, and the first priority value is mapped by a second QoS shaping map into the third priority value, and wherein mapping information for the first QoS shaping map and mapping information for the second QoS shaping map can be

updated during operation of the switch.

23. A switch having one or more switching modules that are capable of QoS shaping and provisioning, each switching module comprising:

one or more ports for receiving a plurality of inbound packets and for transmitting a plurality of outbound packets;
an access controller coupled to the input ports for receiving the inbound packets, each inbound packet having an inbound priority value and a plurality of flow properties; and
a switching controller coupled to the access controller for receiving the inbound packets from the access controller, for determining one or more packet priority values based on the plurality of flow properties, and for providing the outbound packets to the ports to be transmitted.

24. The switch of claim 23 wherein the packet priority values include an internal priority value, and the internal priority value is used to indicate a processing priority of the inbound packets while the inbound packets are being processed in the switch.

25. The switch of claim 23 wherein the packet priority values include an outbound priority value, and the outbound priority value is applied to the inbound packets before the inbound packets are transmitted from the switch as an outbound packet.

26. The switch of claim 23 wherein the packet priority values are determined based on the inbound priority value.

27. The switch of claim 23 wherein the packet priority values are determined based on one or more flow properties including at least one value from a packet field that is not dedicated to defining QoS.

28. The apparatus of claim 23, wherein the packet priority values are determined based on one or more flow properties including at least one of Layer 2 information, Layer 3 information and Layer 4 information.

29. The switch of claim 23 wherein the switching controller includes a priority determination map for mapping the inbound priority value into one or more packet priority values.

30. The switch of claim 29 wherein mapping information of the priority determination map can be updated during operation of the switch.




FIG. 1




FIG. 2




FIG. 3


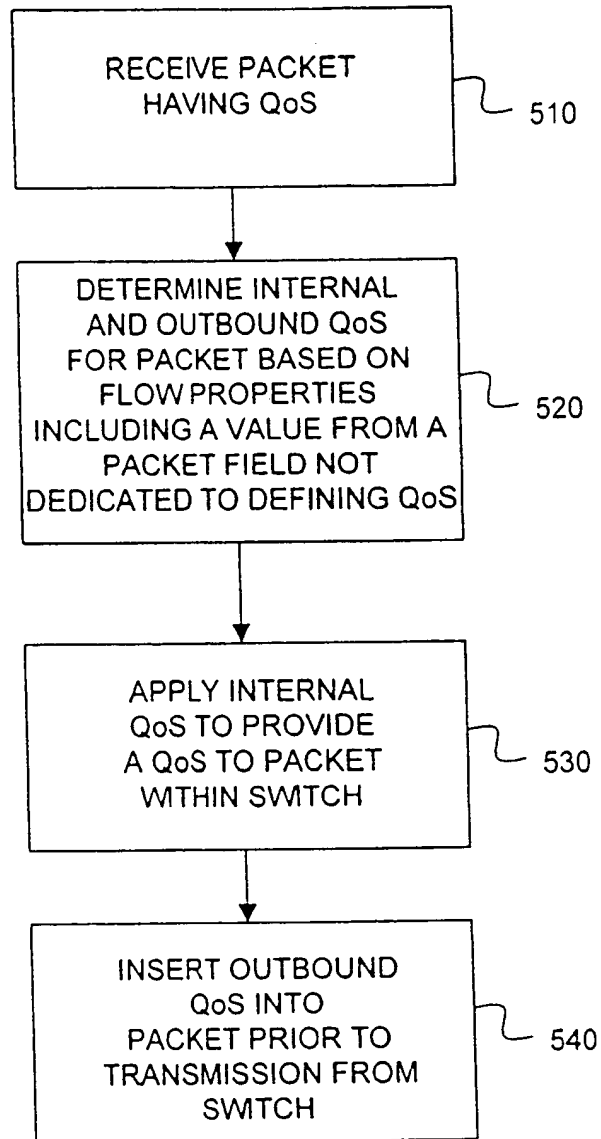
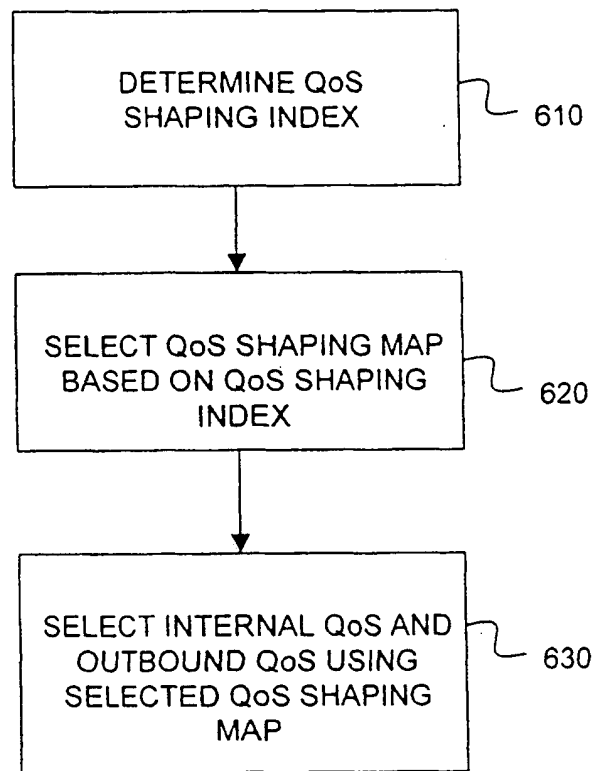


FIG. 4

**FIG. 5**



520

FIG. 6



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 44 0254

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
P,X	WO 00 56024 A (BROADCOM CORP ;KADAMBI SHIRI (US); AMBE SHEKHAR (US); JORDA MICHAEL) 21 September 2000 (2000-09-21) * page 53, line 12 - page 80, line 12 * * page 93, line 5 - page 103, line 4 * * figures 2,14,17 *	1-30	H04L12/56 H04L12/44
X	WO 00 03256 A (BROADCOM CORP) 20 January 2000 (2000-01-20) * page 32, line 3 - page 36, line 11 * * page 46, line 11 - page 52, line 13 * * figures 2,14,17 *	1-30	

TECHNICAL FIELDS
SEARCHED (Int.Cl.7)

H04L

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
BERLIN	18 February 2002	Rosken, W

CATEGORY OF CITED DOCUMENTS
 X particularly relevant (taken alone)
 Y particularly relevant (combined with another document of the same category)
 A technological background
 C non-written disclosure
 P intermediate document

1 theory or principle underlying the invention
 2 earlier patent document, but published only after the filing date
 3 document cited in the application
 4 document cited for other reasons
 5 member of the same patent family, corresponding document

ANNEX TO THE EUROPEAN SEARCH REPORT ON EUROPEAN PATENT APPLICATION NO.

EP 01 44 0254

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

18-02-2002

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 0056024	A	21-09-2000	US 2001043611 A1	22-11-2001
			AU 3529500 A	04-10-2000
			AU 3752000 A	04-10-2000
			AU 4325800 A	04-10-2000
			EP 1161817 A2	12-12-2001
			EP 1169809 A2	09-01-2002
			EP 1159805 A2	05-12-2001
			WO 0056024 A2	21-09-2000
			WO 0056011 A2	21-09-2000
			WO 0056013 A2	21-09-2000
			AU 5442700 A	12-12-2000
			AU 5586400 A	12-12-2000
			AU 5737300 A	31-01-2001
			AU 6334400 A	22-01-2001
			WO 0072533 A1	30-11-2000
			WO 0102965 A2	11-01-2001
			WO 0072531 A1	30-11-2000
			WO 0101724 A2	04-01-2001
WO 0003256	A	20-01-2000	AU 4848499 A	01-02-2000
			AU 4848599 A	01-02-2000
			AU 4848699 A	01-02-2000
			AU 4848799 A	01-02-2000
			AU 6311299 A	01-02-2000
			WO 0003256 A1	20-01-2000
			WO 0003515 A1	20-01-2000
			WO 0003516 A1	20-01-2000
			WO 0003522 A1	20-01-2000
			WO 0003517 A1	20-01-2000
			US 6154446 A	28-11-2000
			US 2001043611 A1	22-11-2001
			US 2001012294 A1	09-08-2001
			US 6104696 A	15-08-2000

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☒ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☒ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

